arXiv:cs/0512091v1 [cs.CG] 23 Dec 2005

Data Structures for Halfplane Proximity Queries and Incremental Voronoi Diagrams*

Boris Aronov† Prosenjit Bose‡ Erik D. Demaine§ Joachim Gudmundsson¶
John Iacono‖ Stefan Langerman** Michiel Smid†

Abstract

We consider preprocessing a set S of n points in the plane that are in convex position into a data structure supporting queries of the following form: given a point q and a directed line ℓ in the plane, report the point of S that is farthest from (or, alternatively, nearest to) the point q subject to being to the left of line ℓ. We present two data structures for this problem. The first data structure uses $O(n^{1+\varepsilon})$ space and preprocessing time, and answers queries in $O(2^{1/\varepsilon} \log n)$ time. The second data structure uses $O(n \log^3 n)$ space and polynomial preprocessing time, and answers queries in $O(\log n)$ time. These are the first solutions to the problem with $O(\log n)$ query time and $o(n^2)$ space.

In the process of developing the second data structure, we develop a new representation of nearest-point and farthest-point Voronoi diagrams of points in convex position. This representation supports insertion of new points in counterclockwise order using only $O(\log n)$ amortized pointer changes, subject to supporting $O(\log n)$-time point-location queries, even though every such update may make $\Theta(n)$ combinatorial changes to the Voronoi diagram. This data structure is the first demonstration that deterministically and incrementally constructed Voronoi diagrams can be maintained in $o(n)$ pointer changes per operation while keeping $O(\log n)$-time point-location queries.

1 Introduction

Line simplification is an important problem in the area of digital cartography [Cro91, Den98, MS92]. Given a polygonal chain P, the goal is to compute a simpler polygonal chain Q that provides a good approximation to P. Many variants of this problem arise depending on how one defines *simpler* and how one defines *good approximation*. Almost all of the known methods of approximation compute distances between P and Q. Therefore, preprocessing P in order to quickly answer distance queries is a common subproblem to most line simplification algorithms.

Of particular relevance to our work is a line simplification algorithm proposed by Daescu et al. [DMSW06]. Given a polygonal chain $P = (p_1, p_2, \ldots, p_n)$, they show how to compute a subchain $P' = (p_{i_1}, p_{i_2}, \ldots, p_{i_m})$, with $i_1 = 1$ and $i_m = n$, such that each segment $[p_{i_j} p_{i_{j+1}}]$ of P' is a good approximation of the subchain of P from p_{i_j} to $p_{i_{j+1}}$. The amount of error is determined by the point of the subchain that is farthest from the line segment $[p_{i_j} p_{i_{j+1}}]$. To compute this approximation efficiently, the key subproblem they solve is the following:

*A preliminary version of this paper appeared in *Proceedings of the 7th Latin American Symposium on Theoretical Informatics*, Valdivia, Chile, March 2006.

†Department of Computer and Information Science, Polytechnic University, Brooklyn, NY, USA. Research supported in part by NSF grant ITR-0081964 and by a grant from US-Israel Binational Science Foundation.

‡School of Computer Science, Carleton University, Ottawa, ON, Canada. Research supported in part by NSERC.

§Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA, USA. Research supported in part by NSF grants CCF-0430849 and OISE-0334653.

¶National ICT Australia, Sydney, Australia.

‖Department of Computer and Information Science, Polytechnic University, Brooklyn, NY, USA. Research supported in part by NSF grants CCF-0430849 and OISE-0334653.

**Chercheur qualifié du FNRS, Département d' Informatique, Université Libre de Bruxelles, Brussels, Belgium.

Problem 1 (Halfplane Farthest-Point Queries). *Preprocess n points $p_1, p_2, \ldots, p_n$ in convex position in the plane into a data structure supporting the following query: given a point q and a directed line ℓ in the plane, report the point p_i that is farthest from q subject to being to the left of line ℓ.*

Daescu et al. [DMSW06] show that, with $O(n \log n)$ preprocessing time and space, these queries can be answered in $O(\log^2 n)$ time. On the other hand, a naïve approach achieves $O(\log n)$ query time by using $O(n^3)$ preprocessing time and $O(n^3)$ space. The open question they posed is whether $O(\log n)$ query time can be obtained with a data structure using subcubic and preferably subquadratic space.

In this paper, we solve this problem with two data structures. The first, relatively simple data structure uses $O(n^{1+\varepsilon})$ preprocessing time and space, and answers queries in $O(2^{1/\varepsilon} \log n)$ time. The second, more sophisticated data structure uses $O(n \log^3 n)$ space and polynomial preprocessing time, and answers queries in $O(\log n)$ time. Both of our data structures apply equally well to halfplane farthest-point queries, described above, as well as the opposite problem of halfplane nearest-point queries. Together we refer to these queries as *halfplane proximity queries*.

Dynamic Voronoi diagrams. An independent contribution of the second data structure is that it provides a new efficient representation for maintaining the nearest-point or farthest-point Voronoi diagram of a dynamic set of points. So far, point location in dynamic planar Voronoi diagrams has proved difficult because the complexity of the changes to the Voronoi diagram or Delaunay triangulation for an insertion can be linear at any one step. The randomized incremental construction avoids this worst-case behavior through randomization. However, for the deterministic insertion of points, the linear worst-case behavior cannot be avoided, even if the points being incrementally added are in convex position, and are added in order (say, counterclockwise). For this specific case, we give a representation of a (nearest-point or farthest-point) Voronoi diagram that supports $O(\log n)$-time point location in the diagram while requiring only $O(\log n)$ amortized pointer changes in the structure for each update. So as not to oversell this result, we note that we do not have an efficient method of determining which pointers to change (it takes $\Theta(n)$ time per change), so the significance of this representation is that it serves as a proof of the existence of an encoding of Voronoi diagrams that can be modified with few changes to the encodings while still supporting point location queries. However, we believe that our combinatorial observations about Voronoi diagrams will help lead to efficient dynamic Voronoi diagrams with fast queries.

Currently, the best incremental data structure supporting nearest-neighbor queries (one interpretation of "dynamic Voronoi diagrams") supports queries and insertions in $O(\log^2 n / \log \log n)$. This result uses techniques for decomposable search problems described by Overmars [Ove83]; see [CT92]. Recently, Chan [Cha06] developed a randomized data structure supporting nearest-neighbor queries in $O(\log^2 n)$ time, insertions in $O(\log^3 n)$ expected amortized time, and deletions in $O(\log^6 n)$ expected amortized time.

2 A Simple Data Structure

In this section, we prove the following theorem:

Theorem 2. *There is a data structure for halfplane proximity queries on a static set of n points in convex position that achieves $O(2^{1/\varepsilon} \log n)$ query time using $O(n^{1+\varepsilon})$ space and preprocessing.*

Our proof is based on starting from the naïve $O(n^3)$-space data structure mentioned in the introduction, and then repeatedly apply a space-reducing transformation. We assume that either all queries are halfplane farthest-point queries or all queries are halfplane nearest-point queries; otherwise, we can simply build two data structures, one for each type of query.

Both the starting data structure and the reduction use Voronoi diagrams as their basic primitive. More precisely, we use the farthest-site Voronoi diagram for the case of halfplane farthest-point queries, and the nearest-site Voronoi diagram for the case of halfplane nearest-point queries. When the points are in convex position and given in counterclockwise order, Aggarwal et al. [AGSS89] showed that either Voronoi diagram can be constructed in linear time. Answering point-location queries in either Voronoi diagram of points in convex position can be done in $O(\log n)$ time using $O(n)$ preprocessing and space [EGS86].

Lemma 3. *There is a static data structure for halfplane proximity queries on a static set of n points in convex position, called* Okey*, that achieves $O(\log n)$ query time using $O(n^3)$ space and preprocessing.*

Proof. Let $p_1, p_2, \ldots, p_n$ denote the n points in convex position in counterclockwise order. The Okey data structure consists of one Voronoi diagram $V(i,j)$ for every subsequence $p_i, p_{i+1}, \ldots, p_j$ of points, where indices are treated modulo n. The space and preprocessing is thus $O(n^3)$.

To answer a halfplane proximity query for a point q and a directed line ℓ, we first find the subsequence of points on the left of line ℓ. In $O(\log n)$ time, we can find the two edges (p_i, p_{i+1}) and (p_j, p_{j+1}) of the convex hull that are intersected by the query line [O'R98, Section 7.9.1]. Then, depending on the orientation of the line ℓ (i.e., which edge is struck first), we can decide between the two possible intervals: $p_{i+1}, p_{i+2}, \ldots, p_j$ or p_{j+1}, p_{j+2}, p_i. Then we locate q in the appropriate Voronoi diagram, either $V(i+1,j)$ or $V(j+1,i)$, and return the site p_k that generated the corresponding Voronoi region. The total query time is $O(\log n)$. □

Transform 4. *Given any static data structure $\mathcal{D}$ for halfplane proximity queries on a static set of n points in convex position that achieves $Q(n)$ query time using $M(n)$ space and preprocessing, and for any parameter $m \le n$, there is a static data structure for halfplane proximity queries on a static set of n points in convex position, called $\mathcal{D}$-Dokey, that achieves $2\,Q(n) + O(\log n)$ query time using $\lceil n/m \rceil\, M(m) + O(n^2/m)$ space and preprocessing.*

Proof. Let $p_1, p_2, \ldots, p_n$ be the n points in convex position in counterclockwise order. We define the $\lceil n/m \rceil$ *breakpoints* to be the points p_k with $k \equiv 1 \pmod m$, i.e., the points p_{im+1} for $i \in \{0, 1, \ldots, \lceil n/m \rceil - 1\}$. The data structure consists of two substructures:

$\mathcal{D}$ Substructure: We construct an instance of the data structure $\mathcal{D}$ on the half-open interval of points between every consecutive pair of breakpoints. More precisely, for each $i \in \{0, 1, \ldots, \lceil n/m \rceil - 1\}$, we construct an instance of $\mathcal{D}$ on the points $p_{im+1}, p_{im+2}, \ldots, p_{\min\{n,(i+1)m\}}$. These structures require $\lceil n/m \rceil (M(m) + O(1))$ space and preprocessing.

Voronoi Substructure: For each breakpoint p_k, we construct Voronoi diagrams on all intervals of points of length an exact power of two with one endpoint at p_k. More precisely, for each $i \in \{0, 1, \ldots, \lceil n/m \rceil - 1\}$, and for each $j \in \{0, 1, \ldots, \lfloor \log n \rfloor\}$, we construct the Voronoi diagram on the points p_{im+1+r} for $r \in \{0, 1, \ldots, 2^j - 1\}$, and we construct the Voronoi diagram on the points p_{im+1-r} for $r \in \{0, 1, \ldots, 2^j - 1\}$, where indices are treated modulo n. The space and preprocessing requirements for these Voronoi diagrams are

$$O\left(|S_m| \sum_{j=0}^{\lfloor \log n \rfloor} 2^j\right) = O\left(\frac{n}{m} \cdot n\right) = O\left(\frac{n^2}{m}\right).$$

Overall, the space and preprocessing required for $\mathcal{D}$-Dokey is $\lceil n/m \rceil\, M(m) + O(n^2/m)$ as claimed.

It remains to show how we can use $\mathcal{D}$-Dokey to answer halfplane proximity queries in $2\,Q(n) + O(\log n)$ time. Suppose that we are given a point q and a directed line ℓ. As described in the proof of Lemma 3, in $O(\log n)$ time, we can find the interval $p_i, p_{i+1}, \ldots, p_j$ of points to the left of line ℓ. If this interval contains no breakpoints, then it is contained in the interval of a $\mathcal{D}$ substructure, so we can answer the query in $Q(n)$ time by asking the $\mathcal{D}$ substructure the same halfplane proximity query. Otherwise, let $p_{i'}$ and $p_{j'}$ be the first and the last breakpoints in the interval, respectively. We ask the $\mathcal{D}$ substructure immediately preceding $p_{i'}$ (representing the interval $p_{i'-m}, p_{i'-m+1}, \ldots, p_{i'-1}$ if $i' > 0$, or else the interval $p_{(\lceil n/m \rceil - 1)m+1}, p_{(\lceil n/m \rceil - 1)m+2}, \ldots, p_n$) and the $\mathcal{D}$ substructure immediately succeeding $p_{j'}$ (representing the interval $p_{j'}, p_{j'+1}, \ldots, p_{\min\{n, j'+m-1\}}$) the same halfplane proximity query. These queries cover the ranges $p_i, p_{i+1}, \ldots, p_{i'-1}$ and $p_{j'}, p_{j'+1}, \ldots, p_j$. To cover the remaining range $p_{i'}, p_{i'+1}, \ldots, p_{j'}$ between the two breakpoints, we use the property that any interval can be covered (with overlap) by two intervals of length an exact power of two. Namely, let $k = 2^{\lfloor \lg(j'-i') \rfloor}$, where the difference $j' - i'$ accounts for wraparound modulo n. We query q in the Voronoi diagram on the interval $p_{i'}, p_{i'+1}, \ldots, p_{i'+k}$ and in the Voronoi diagram on the interval $p_{j'}, p_{j'-1}, \ldots, p_{j'-k}$. Together, the four queries cover (with overlap) the desired interval $p_i, p_{i+1}, \ldots, p_j$. Among the four results from the four queries, we return the best (either farthest or nearest) relative to point q. □

By starting with the data structure Okey of Lemma 3, and repeatedly applying the Dokey transformation of Transformation 4, we obtain the structure Okey-Dokey-Dokey-Dokey-..., or Okey-Dokey^k, which leads to the following:

Corollary 5. *For every integer $k \geq 1$, Okey-Dokey^{k-1} is a data structure for halfplane proximity queries on a static set of n points in convex position that achieves $O(2^k \log n)$ query time using $O(n^{(2k+1)/(2k-1)})$ space and preprocessing.*

Proof. The proof is by induction on k. In the base case $k = 1$, we can use the Okey data structure from Lemma 3 because $(2k+1)/(2k-1) = 3$. For $k > 1$, assume by induction that we have a data structure that achieves query time at most $c\,(2^{k-1}-1)\log n$ using space and preprocessing at most $c\,n^{(2k-1)/(2k-3)}$. Assume that the constant c is at least twice as large as the constants implicit in the O notation in Transform 4. We apply the Dokey transformation from Transform 4 to this data structure, substituting $m = n^{(2k-3)/(2k-1)}$. Thus, $n/m = 1/n^{2/(2k-1)}$ and $n^2/m = n^{(2k+1)/(2k-1)}$. The resulting query time is at most $2\,c\,(2^{k-1} - 1)\log n + (c/2)\log n \leq c\,(2^k - 1)\log n$ as desired. The resulting space and preprocessing time is at most $(n/m + 1)\,c\,m^{(2k-1)/(2k-3)} + (c/2)\,n^2/m = c\,(1/n^{2/(2k-1)} + 1)n + (c/2)\,n^{(2k+1)/(2k-1)} = c\,n^{(2k-3)/(2k-1)} + c\,n + (c/2)\,n^{(2k+1)/(2k-1)} \leq c\,n^{(2k+1)/(2k-1)}$ for sufficiently large n, as desired. □

The space and preprocessing time of Okey-Dokey^{k-1} according to Corollary 5 can be written as $n^{1+2/(2k-1)}$. For any given $\varepsilon > 0$, we choose $k = 1/2 + 1/\varepsilon$. Then the space and preprocessing time are $O(n^{1+\varepsilon})$ and the query time is $O(2^{1/\varepsilon}\log n)$, proving Theorem 2.

3 Grappa Trees

Our faster data structure for halfplane proximity queries requires the manipulation of binary trees with a fixed topology determined by a Voronoi diagram. To support efficient manipulation of such trees, we introduce a data structure called *grappa trees*. This data structure is a modification of Sleator and Tarjan's link-cut trees [ST83] that supports some unusual additional operations.

Definition 6. *Grappa trees solve the following data-structural problem: maintain a forest of rooted binary trees with specified topology subject to*

T = Make-Tree(v): *Create a new tree T with a single vertex v (not previously in another tree).*

T = Link(v, w, d, m_ℓ, m_r): *Given a vertex v in some tree T_v and the root w of a different tree T_w, add an edge (v, w) to make w a child of v, merging T_v and T_w into a new tree T. The value $d \in \{\ell, r\}$ specifies whether w becomes a left or a right child of v; such a child should not have existed previously. The new edge (v, w) is assigned a* left mark *of m_ℓ and a* right mark *of m_r.*

(T_1, T_2) = Cut(v, w): *Delete the existing edge (v, w), causing the tree T containing it to split into two trees, T_1 and T_2. Here one endpoint of (v, w) becomes the root of the tree T_i that does not contain the root of T.*

Mark-Right-Spine(T, m): *Set the right mark of every edge on the right spine of tree T (i.e., the edge from the root of T to its right child, and recursively such edges in the right subtree of T) to the new mark m, overwriting the previous right marks of these edges.*

(e, m_ℓ^*, m_r^*) = Oracle-Search(T, O_e): *Search for the edge e in tree T. The data structure can find e only via* oracle queries*: given two incident edges (u, v) and (v, w) in T, the oracle $O_e(u, v, w, m_\ell, m_r, m'_\ell, m'_r)$ determines in constant time which of the subtrees of $T - v$ contains x.*[1] *(Note that edges (u, v) and (v, w) are considered to exist in $T - v$, even though one of their endpoints has been removed.) The data structure provides the oracle with the left mark m_ℓ and the right mark m_r of (u, v), as well as the left mark m'_ℓ and the right mark m'_r of (v, w), and at the end, it returns the left mark m_ℓ^* and the right mark m_r^* of the found edge e.*

[1] Given the number of arguments, it is tempting to refer to the oracle as $O(A, B, D, G, I, L, S)$, but we will resist that temptation.

Theorem 7. *There exists an $O(n)$-space constant-in-degree pointer-machine data structure that maintains a forest of grappa trees and supports each operation in $O(\log n)$ worst-case time per operation, where n is the total size of the trees affected by the operation.*

Proof. Our grappa-tree data structure is based on the worst-case version of the link-cut tree data structure of Sleator and Tarjan [ST83, Section 5]. This data structure maintains a forest of fixed-topology trees subject to Make-Tree, Link, Cut, and several other operations, each in $O(\log n)$ worst-case time per operation, and using $O(n)$ space. The data structure represents each tree in the forest by decomposing it into a set of disjoint paths connected by tree edges. Each path is in turn represented by a biased binary tree whose nodes represent the vertices of the path, ordered in the biased tree according to the depth along the path. Thus, higher vertices in the path correspond to nodes farther left in the biased tree. (In fact, the structure in [ST83] puts the vertices of the path at the leaf nodes of the biased tree and uses the internal nodes of the biased tree to represent the edges of path, but it is easy to modify to use this view.) Thus, there is a bijection between edges of a path and edges in the biased tree representing that path. The link-cut tree structure for a binary tree T can therefore be seen as a tree R, the *representation tree*, in which each node has at most four children (up to two for the biased tree representing a path and up to two for interconnecting the paths/biased trees), the nodes of R correspond bijectively to vertices of T, and the edges of R correspond bijectively to edges of T. Note that the endpoints of an edge in R may differ from the endpoints of the corresponding edge in T, although the parent endpoint of the edge in R always corresponds to one of the endpoints of the corresponding edge in T (which endpoint depends on whether the edge is a left edge or a right edge in R). For an edge (v, w) between parent v and child w in R, if we restrict to the nodes in the subtree of R rooted at v, then the restricted sets of nodes on either side of (v, w) in R correspond exactly to the restricted sets of vertices on either side of the corresponding edge in T. Furthermore, by suitable biasing as described in [ST83], R has height $O(\log n)$.

We augment the representation tree R to enable marking as follows. Because our tree T has bounded degree (an assumption not made in [ST83]), we can also explicitly store T (the parent, left child, and right child of each vertex) and crosslink between corresponding nodes/vertices and corresponding edges in the two structures. To each edge of T we add a left-mark field and a right-mark field. These fields contain the last explicitly stored marks for the edge, and for edges connecting two different paths (*nonpath edges*), they are accurate, while for edges on a path, the right-mark field may become out-of-date. In particular, Link(v, w, m_ℓ, m_r) sets the left-mark and right-mark fields of the created edge (v, w) to the specified values m_ℓ and m_r, respectively. To each internal node of a biased tree in R, we also add a right-mark field, which may be blank. When nonblank, this field represents bulk right markings that should be (but have not yet been) applied to the entire subtree of the biased tree rooted at this node. Thus, the actual right mark of an edge e on a path in T is implicitly the first nonblank right-mark field of a node along the path from the root of the biased tree representing the path to the corresponding edge of e in R, if there is such a field, or else the right-mark field of the edge e itself.

We can maintain this augmentation as the representation tree R changes. Because the definition of the augmented values is relative to individual biased trees, we care only about modifications to biased trees themselves, not about the modifications to the edges between different biased trees that form the entire representation tree R. The link-cut data structure modifies biased trees according to rotations, splits, and concatenations. We can modify the implementation of all of these operations to propagate the mark fields, at the cost of an extra constant factor, in such a way that preserves the implicit marks of all edges in T. The idea is to push down node marks judiciously: whenever any operation visits a node v in R with a nonblank right-mark field, copy that value to the right-mark fields of the edges in T corresponding to the up to two children edges of v in the biased tree, copy the value to the right-mark fields of the up to two children nodes of v in the biased tree, and then blank out the field in the node v itself. Because operations on link-cut trees always start at the root of R and traverse along paths down from there, any nodes involved in the operation will have already cleared their mark fields before they actually get used, so the marks on the corresponding edges in T will be up-to-date.

To implement Mark-Right-Spine(T, m), we visit all biased trees that represent paths with edges along the right spine of T. We start at the biased tree at the root of R (which contains the node corresponding to the root vertex of T), and we walk to the leftmost edge in the biased tree that corresponds to a left edge in T. (This edge can be found by a simple augmentation to store which subtrees in a biased tree contain edges

corresponding to left edges in T.) The edges that precede e in the biased tree correspond to the portion of the right spine of T on the corresponding path. Some of these edges preceding e are on or adjacent to the path we just walked, and we set the right-mark fields of the corresponding right-spine edges in T explicitly. The remainder of the edges preceding e are in subtrees hanging off the left of the path we just walked, and we set the right-mark fields of the root nodes of these subtrees. The total number of markings is at most proportional to the length of our walk. Then we continue to the child biased tree that precedes e, if any, which represents the continuation of the right spine of T. (Again this child biased tree can be found by a simple augmentation to store which subtrees in a biased tree have a descendant child biased tree.) Because R has height $O(\log n)$, the entire length of the walk and thus the total number of markings is $O(\log n)$.

Given a query oracle O_e and a tree T, we can perform Oracle-Search by a tree walk in R starting at the root. Upon visiting a node v with two children, we apply the oracle to the two children edges of v, which correspond to two incident edges of T sharing the vertex corresponding to v. The result of this oracle call tells us whether e is the parent edge of v or else which of the two children subtrees contains e. Here we use the correspondence between the sides of an edge in R and the sides of the corresponding edge in T, assuming that we have already narrowed our search to the subtree of R rooted at v. On the other hand, if we visit a node v with just one child, we apply the oracle to the unique child edge and the parent edge of v, which again correspond to two incident edges of T sharing the vertex corresponding to v. This oracle call tells us whether e is the parent edge of v or e is in the child subtree. Because R has height $O(\log n)$, Oracle-Search queries run in $O(\log n)$ worst-case time. Furthermore, any edge we visit during this traversal from the root will have its left-mark field and right-mark field up-to-date, because the ancestor nodes above the edge will have already been visited and thus their fields will have already been propagated and blanked, so we can provide the oracle with the left and right marks of each query edge. Similarly, once we locate the edge e, we know its left and right labels. □

4 Rightification of a Tree: Flarbs

The fixed-topology binary search tree maintained by our faster data structure for halfplane proximity queries changes in a particular way as we add sites to a Voronoi diagram. We delay the specific connection for now, and instead define the way in which the tree changes: a tree restructuring operation called a "flarb". Then we bound the work required to implement a sequence of n flarbs by showing that the total number of pointers changes (i.e., the total number of parent/left-child and parent/right-child relationships that change) is $O(n \log n)$. Thus, for the remainder of this section, we use the term *cost* to refer to (a constant factor times) the number of pointer changes required to implement a tree-restructuring operation, not the actual running time of the implementation. This bound on cost will enable us to implement a sequence of n flarbs via $O(n \log n)$ link and cut operations, for a total of $O(n \log^2 n)$ time.

The flarb operation is parameterized by an "anchored subtree" which it transforms into a "rightmost path". An *anchored subtree* S of a binary search tree T is a connected subgraph S of T that includes the root of T. A *right-leaning path* in a binary search tree T is a path monotonically descending through the tree levels, always proceeding from a node to its right child. A *rightmost path* in T is a right-leaning path that starts at the root of T.

The *flarb* operation[2] of an anchored subtree S of a binary search tree T is a transformation of T defined as follows; refer to Figure 1. First, we create a new root node r with no right child and whose left child subtree is the previous instance of T; call the resulting binary search tree T'. We extend the anchored subtree S of T to an anchored subtree S' of T' by adding r to S. Now we re-arrange S' into a rightmost path on the same set of nodes, while maintaining the binary search tree order (in-order traversal) of all nodes. The resulting binary search tree T'' is the result of flarbing S in T.

Theorem 8. *A sequence of n flarb operations, starting from an empty tree, can be implemented at a cost of $O(\log n)$ amortized pointer changes per flarb.*

[2] "Flarb" is a clever abbreviation of a long technical term whose meaning we cannot reveal for reasons we cannot comment on at the moment, perhaps simply due to lack of space or of the aforementioned purported meaning. Note that this notion of flarb is different from that of [Cal05].

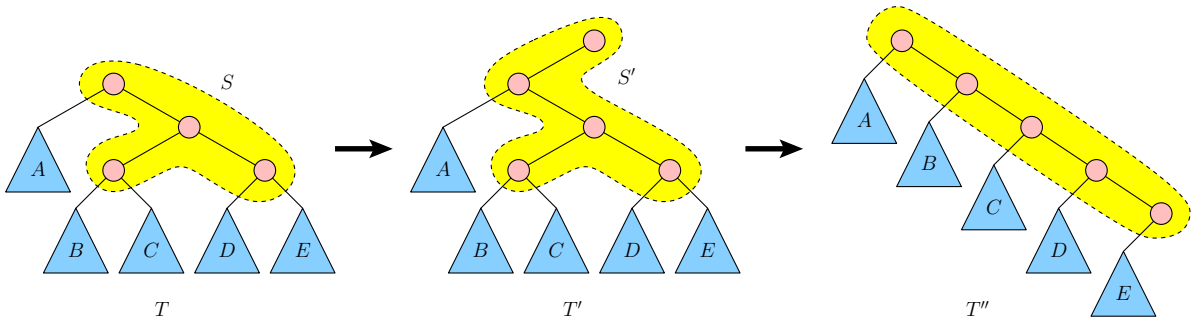

Figure 1: An example of a flarb. The anchored subtree is highlighted.

Proof. We use the potential method of amortized analysis, with a potential function inspired by the analysis of splay trees [ST85]. For any node x in a tree T, let $w(x)$ be the *modified weight* of the subtree rooted at x, which is the number of nodes in the subtree plus the number of null pointers in the tree. In other words, we add dummy nodes as leaves in place of each null pointer in T, for the purpose of computing subtree size. Define $\varphi(x) = \lg \frac{w(left(x))}{w(right(x))}$. Clearly $|\varphi(x)| \leq \lg(2n-1)$, because the smallest possible subtree contains no real nodes and one dummy node, and the largest possible subtree contains $n-1$ real nodes and n dummy nodes. The potential of a tree T with n nodes is $\Phi(T) = \sum_x \varphi(x)$, with the sum taken over the (real) nodes x in T. Therefore, $|\Phi(T)| = O(n \log n)$ for any tree T.

For the purposes of the analysis, we use the following heavy-path decomposition of the tree. The *heavy path* from a node continues recursively to its child with the largest subtree, and the *heavy-path decomposition* is the natural decomposition of the tree into maximal heavy paths. Edges on heavy paths are called *heavy edges*, while all other edges (connecting two heavy paths) are called *light edges*.

Outline. To analyze a flarb in a binary search tree T, we decompose the transformation into a sequence of several steps, and analyze each step separately.

First, the addition of the new root node r can be performed by changing a constant number of pointers in the tree. Because $\varphi(r) = \lg(2n-1)$, the amortized cost of this operation is trivially $O(\log n)$. Thus, in the remainder of the proof, we focus on the actual restructuring of the resulting anchored subtree S' into a rightmost path, a process we call *rightification*.

At all times during rightification, the nodes constituting the original anchored subtree S' continue to form an anchored subtree of the current binary search tree, and for simplicity of notation we continue to denote the current such anchored subtree as S'.

To implement rightification, we first execute several simplifying steps of two types, called "zig" and "zag",[3] in no particular order. Each such step has zero amortized cost. Any number of such operations might need to be performed and we stop when neither can be applied. At this point, the anchored subtree S' has a particular form and we perform a final operation, called a "stretch", at the cost of $O(\log n)$ amortized pointer changes. This bound, together with the observation that the potential drop over any sequence of operations is $O(n \log n)$, gives the theorem. We now describe the details of zig-zagging and stretching.

The zig. A *zig* is executed whenever a light left edge is part of the anchored subtree S'; see Figure 2. The zig operation simply involves a right rotation on the edge in question. The actual cost of a zig is $O(1)$, which we set to be 1 to ease the analysis.

To analyze the change in potential, let A and B denote the two children subtrees of the lower endpoint of the edge, and let C denote the right child subtree of the upper endpoint of the edge. We use the same letters to denote the modified weight of the subtree. Because the edge is light, $A + B + 1 \leq C$. Then the

[3]Unlike most terminology in this paper, these terms are used for no particular reason. Cf. footnote 2.

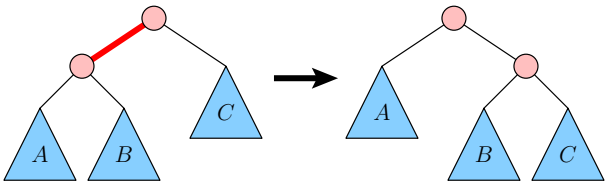

Figure 2: A zig: The thick edge belongs to the anchored subtree S' and is light.

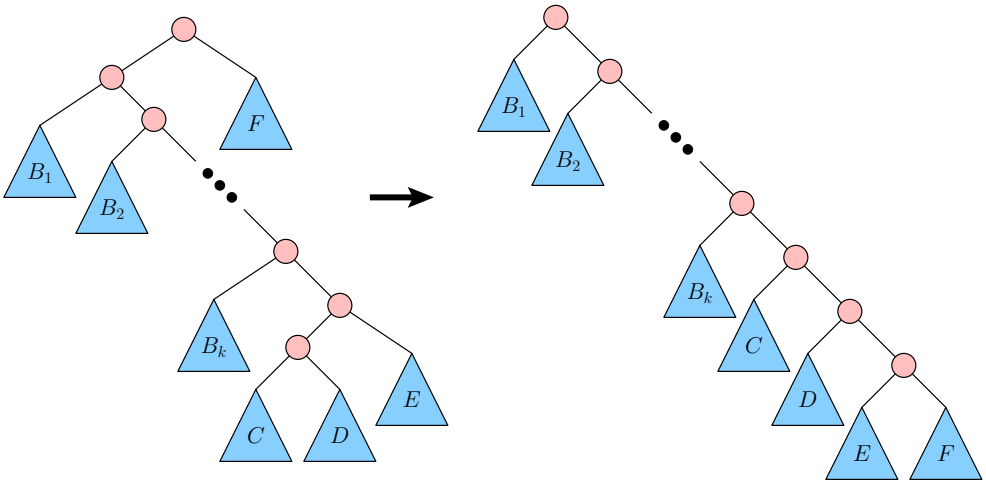

Figure 3: A zag.

potential change is

$$\begin{aligned}\Delta\Phi &= \lg\frac{A}{B+C+1} + \lg\frac{B}{C} - \lg\frac{A}{B} - \lg\frac{A+B+1}{C} \\ &= \lg\underbrace{\frac{B}{A+B+1}}_{<1} + \lg\underbrace{\frac{B}{B+C+1}}_{<\frac{1}{2}\text{ because }C>B} < -1.\end{aligned}$$

Therefore, the amortized cost of a zig is zero, as claimed.

The zag. A *zag* is performed whenever there exists, within the anchored subtree S', a path that goes left one edge, right zero or more edges, and then left again one edge; see Figure 3. The zag operation performs a constant number of pointer changes to re-arrange the path in question into a right-leaning path. The actual cost of a zag is $O(1)$, which we again set to be 1 to ease the analysis.

We now argue that a zag reduces the potential by at least 1. First, notice that the contribution to the potential of parent nodes of the trees $B_1, B_2, \ldots, B_k$ decreases after the execution of the zag because, in each case, the left subtree remains the same while the right subtree grows. We will argue that the contribution

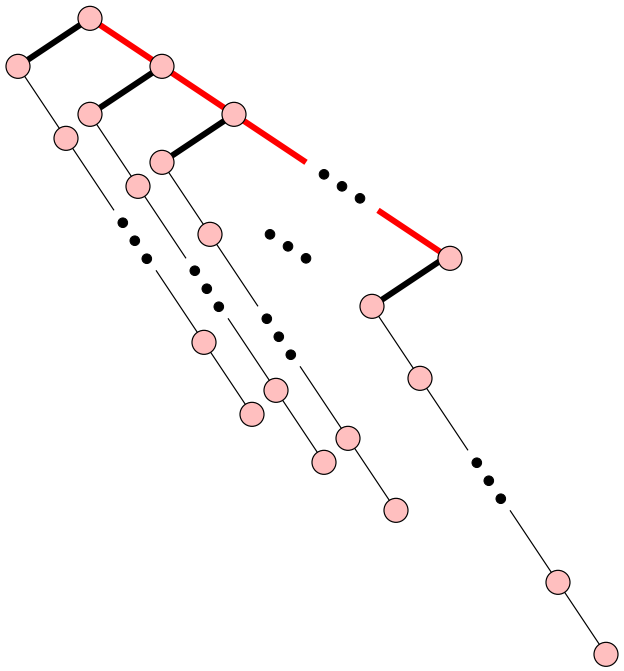

Figure 4: The form of the anchored subtree S' before the final stretch. The thick light edges are light, and the thick black edges are heavy.

of the remaining nodes decreases by at least 1. Indeed,

$$\begin{aligned}
\Delta\Phi &\le \lg \frac{C}{D+E+F+2} + \lg \frac{D}{E+F+1} + \lg \frac{E}{F} \\
&\quad - \lg \frac{\sum_i B_i + C + D + E + k + 2}{F} - \lg \frac{C+D+1}{E} - \lg \frac{C}{D} \\
&= \lg \frac{D^2 E^2}{\underbrace{(D+E+F+2)}_{>D+E}(E+F+1)\underbrace{(\Sigma_i B_i + C + D + E + k + 2)}_{>D+E}(C+D+1)} \\
&< \lg \Big(\frac{1}{2} \cdot \underbrace{\frac{2DE}{(D+E)^2}}_{<1} \cdot \underbrace{\frac{D}{C+D+1}}_{<1} \cdot \underbrace{\frac{E}{E+F+1}}_{<1} \Big) < -1,
\end{aligned}$$

as claimed.

The final stretch. After all possible zigs and zags have been exhausted, we claim that the anchored subtree S' must have the form shown in Figure 4. Indeed, any tree that has no light left edge and no right-leaning path delimited by two left edges must have this form. In particular, because the rightmost path in this tree must be light, its length is at most $\lg(2n+1)$.

The final *stretch* operation, which completes the flarb, simply converts this tree into a rightmost path by effectively concatenating the subsidiary right-leaning paths, incorporating them into the main path. Only $O(\log n)$ actual pointer changes are required. The potential does not increase because left subtrees of every node shrink and right subtrees grow, if they change at all. Therefore, the amortized cost of the stretch is indeed $O(\log n)$.

This concludes the proof of the theorem. □

5 Transformations

In this section we show how flarbs and grappa trees come together with a little work to give us the main result. In the next two transformations, we focus on the farthest-point case, but the proofs apply equally well to nearest-point.

Transform 9. *Given a grappa tree data structure supporting each operation in $O(\log n)$ worst-case time, and given a data structure to incrementally maintain a tree created by n flarbs with $O(\log n)$ amortized pointer changes per flarb, we can construct an $O(n \log^2 n)$-space data structure that supports $O(\log n)$-time farthest-point queries on any prefix of a sequence of points in convex position in counterclockwise order.*

Proof. We construct an incremental data structure that supports $O(\log n)$-time farthest-point queries on the current sequence of points, $\langle p_1, p_2, \ldots, p_n \rangle$, and supports appending a new point p_{n+1} to the sequence provided that this change maintains the invariant that the vertices remain in convex position and in counterclockwise order. Thus the insertion order equals the index order and equals the counterclockwise traversal order of a convex polygon. The data structure runs on a pointer machine in which each node has bounded in-degree. Thus we can apply the partial-persistence transform of [DSST89] and obtain the ability to support farthest-point queries on any prefix of the inserted points in $O(\log n)$ time. The space usage becomes proportional to the number of pointer changes during the insertions.

We consider the ordered tree T formed by the finite segments of the farthest-point Voronoi diagram, ignoring their precise geometry; see Figure 5. More precisely, the *farthest-point Voronoi diagram* [PS93, Section 6.3] divides the plane into n cells by classifying each point q in the plane according to which of $p_1, p_2, \ldots, p_n$ is the farthest from q. The *farthest-point Delaunay triangulation* [Epp92] is the dual of the farthest-point Voronoi diagram, i.e., it triangulates the convex polygon with vertices $p_1, p_2, \ldots, p_n$ by connecting two vertices whenever the corresponding Voronoi cells share an edge. We consider the dual tree T of this farthest-point Delaunay triangulation of the convex polygon, i.e., the dual graph excluding the infinite region exterior to the convex polygon. Each edge in this tree corresponds to (a nongeometric representation of) a finite edge of the farthest-point Voronoi diagram, which is the bisector of two of the points p_i and p_j that are adjacent in the Delaunay triangulation. Each node in the tree represents a vertex in the farthest-point Voronoi diagram, or equivalently a triangle in the farthest-point Delaunay triangulation, and therefore has degree $d \leq 3$, where any degree deficit corresponds to $3 - d$ infinite rays in the farthest-point Voronoi diagram not represented in the tree T.

We can view the tree T as a binary search tree as follows. First, we root the tree at the node corresponding to the unique triangle in the Delaunay triangulation bounded by the edge connecting the first inserted point p_1 and the most recently inserted point p_n. We view the infinite ray emanating from the Voronoi vertex as the "parent edge" of this root node, defining the notion of *left child* versus *right child* of a node according to the counterclockwise order around the Voronoi vertex. (Note that this order is the opposite of the order defined by the triangulation, so in Figure 5 (right), we draw T in mirror image so that its geometric notions of left and right match that of the Voronoi diagram.) Second, we assign keys to nodes consistent with the in-order traversal. For each tree node corresponding to a Delaunay triangle with vertices p_i, p_j, p_k, where $i < j < k$, we assign a key of j. In other words, we assign the median of the three vertex labels of the Delaunay triangle to be the key of the corresponding tree node.

One way to view this key assignment is as follows. If we imagine adding an infinite rays in place of each absent child in the tree, and add an infinite ray in place of the absent parent of the root (the dashed lines in Figure 5, right), matching the counterclockwise order around the Voronoi vertex, then we decompose the plane into regions corresponding to Voronoi regions, each of which corresponds to a single point p_i. All of the nodes bounding p_i's region correspond to triangles incident to p_i. We assign the key i to the unique such node in T that is closest to the root of T, or equivalently the least common ancestor of such nodes, which is the inflection point between two descendant paths that bound the region. Two exceptions are $i = 1$ and $i = n$: the vertices incident to p_1 are those on the left spine of T, and the vertices incident to p_n are those on the right spine of T.

In this view, we also define the *left mark* of an edge to be the label of the region to the left of the edge, and similarly for the *right mark*. Thus, the two marks of an edge define the two points p_i and p_j whose bisector line contains the Voronoi edge. If an edge is the left edge of its parent node, then the edge's right mark is simply the key of that parent, because the right edge of the parent creates an inflection point at the

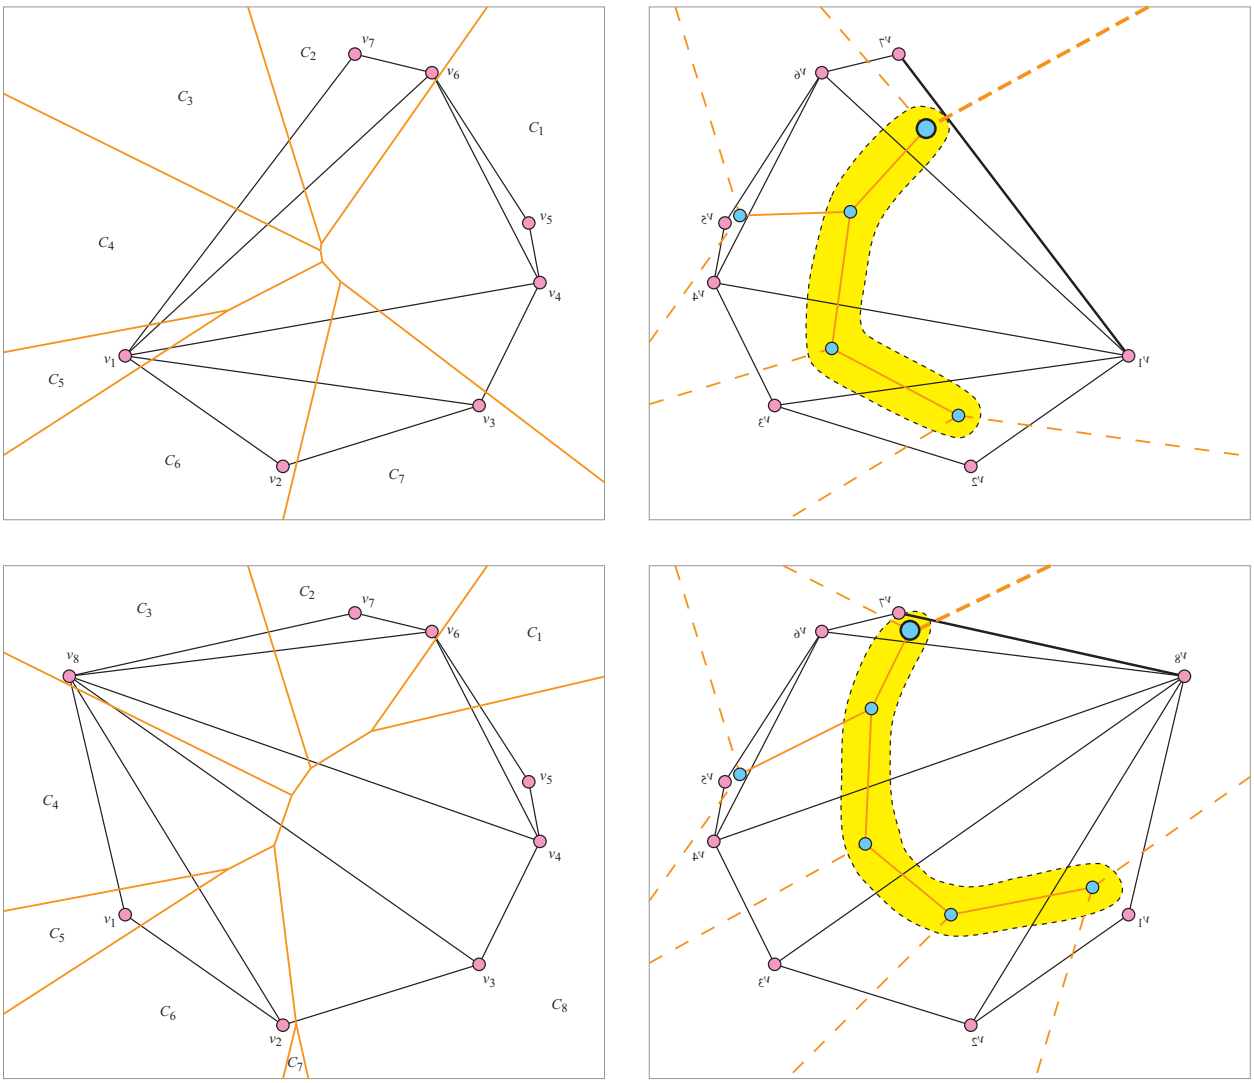

Figure 5: Adding vertex v_8 in counterclockwise order. Top: Before. Bottom: After. Left: Farthest-point Voronoi diagram and its dual, the Delaunay triangulation. Right: Delaunay triangulation and its dual, the tree T with attached infinite rays drawn as dashed lines, drawn in mirror image so that geometric left versus right matches the order in the Voronoi diagram. The root vertex of T and its parent edge are emboldened.

parent. Similarly, if an edge is the right edge of its parent node, then the edge's left mark is the key of that parent. Intuitively, in either case, if we walk up from the edge on its "underside", then we immediately find a local maximum in the region. On the other hand, in either case, the other mark of the edge is the key of the parent node of the deepest ancestor edge that has the opposite orientation (left versus right): this bending point is the first inflection point we encounter as we walk up the tree on the "top side" of the edge. The tree T is not balanced, so we use a grappa tree to represent it and the left and right marks of edges.

Next we consider the effect of inserting a new point p_{n+1}. As in the standard incremental algorithm for Delaunay construction [dBvKOS99, Section 9.3], we view the changes to the farthest-point Delaunay triangulation as first adding a triangle p_1, p_n, p_{n+1} and then flipping a sequence of edges to restore the farthest-point Delaunay property. The key property of the edge-flipping process is that all flipped edges end up incident to the newly inserted point p_{n+1}. Therefore these changes can be interpreted in the tree as adding a new root node, whose left child is the previous root, and then choosing a collection of nodes to move to the right path of the new root. This collection of nodes induces a connected subtree because the triangles

involved in the flips form a connected set. (In particular, the flipping algorithm considers the neighbors of a triangle for flipping only if the triangle was already involved in a flip.) Thus, the changes correspond exactly to a flarb, with the flexibility of the flarb operation encompassing the various possibilities of which edges get flipped to maintain the farthest-point Delaunay property. Another way to view the addition of p_{n+1} is directly in the Voronoi diagram. The point p_{n+1} will capture the region R_{n+1} for which p_{n+1} is the farthest neighbor. The region R_{n+1} is a convex polygon. Outside R_{n+1}, the Voronoi diagram is unchanged, so all edges of the new Voronoi diagram are either bisectors of the same two points as before, or are edges of R_{n+1}. In T after the flarb, R_{n+1} corresponds to the right spine.

Each pointer change during a flarb operation can be implemented with one cut and one link operation. Therefore the grappa tree implements the $O(n \log n)$ total pointer updates from flarb operations in $O(n \log^2 n)$ total pointer updates. It remains to update the marks on the edges. By the incremental Voronoi/Delaunay view above, the only edges for which these marks might change are the edges incident to the new region R_{n+1}, i.e., the edges on the right spine. We update the right marks on all of these edges by calling Mark-Right-Spine$(T, n+1)$. The left mark of each edge on the right spine is simply the key of the parent node of the edge. During the execution of the flarb, various right paths were cut and pasted together with cuts and links to form the final right spine. The edges on the final right spine that were originally part of a right path in T already had a left mark equal to the key of their parent node. Any other edges on the final right spine were just added via links, so their left marks can be set accordingly by specifying the right m_ℓ argument to Link. Thus, the total number of pointer updates remains $O(n \log^2 n)$. This concludes the space bound of the data structure.

To support farthest-point queries, it suffices to build an oracle for the grappa tree's Oracle-Search. Specifically, given two incident edges (u, v) and (v, w), the oracle must determine which subtree of $T - v$ has the answer to the farthest-point query. Using the two marks on the two edges, two of which must be identical, we can determine the three vertices p_i, p_j, and p_k of the Delaunay triangle corresponding to vertex v in T. The vertex of the Voronoi diagram corresponding to v lies at the intersection of the three perpendicular bisectors between these three vertices of the Delaunay triangle. We draw three rays from this Voronoi vertex to each of the three corners of the Delaunay triangle. These three rays divide the plane into three sectors, and the Voronoi regions corresponding to the nodes in each subtree of $T - v$ lie entirely in one of these sectors, with exactly one subtree per sector. In constant time, we can decide which of the three sectors contains the query point q. The farthest-point Voronoi region containing the query point q is guaranteed to be incident to the corresponding subtree, and therefore we obtain a suitable answer for the oracle query. At the end, Oracle-Search will narrow the search to a specific edge of T, meaning that the query point q is in one of the two Voronoi regions incident to the corresponding Voronoi edge. In constant time, using the two labels on that edge of the tree, we can determine which side of the bisector contains q, and therefore which Voronoi region contains q, i.e., which point p_i is farthest from q.

This concludes the proof of the theorem. □

Transform 10. *Given an $O(n \log^2 n)$-space data structure that supports $O(\log n)$-time farthest-point queries on any prefix of a sequence of n points ordered in convex position in counterclockwise order, we can construct an $O(n \log^3 n)$-space data structure that supports $O(\log n)$-time farthest-point-left-of-line queries on n points in convex position.*

Proof. Let $p_1, p_2, \ldots, p_n$ denote the n points in counterclockwise order.

First we observe that, using the given prefix structure, we can also build an $O(n \log^2 n)$-space data structure that supports $O(\log n)$-time farthest-point queries on any suffix of a sequence of n points ordered in convex position in counterclockwise order. We simply reflect the points about a fixed axis, reverse the order of the points, and build the prefix structure, and then apply the same reflection transformation to query points before giving it to the structure.

Next we observe that, in $O(\log n)$ time, we can find the interval $p_i, p_{i+1}, \ldots, p_j$ of points that are to the left of the query line. This algorithm is described in the proof of Lemma 3.

We build a collection of prefix and suffix data structures, and answer a query, via a divide-and-conquer recursion. The top level of the recursion is special because the sequence $p_1, p_2, \ldots, p_n$ is cyclic. In this case we build a prefix structure and a suffix structure on this list of points. These structures can be used to solve any query interval that contains either p_1 or p_n or both. Namely, if interval contains exactly one of p_1 or p_n,

then the interval is a prefix or suffix of $p_1, p_2, \ldots, p_n$. Otherwise, the interval is the union of a prefix and a suffix. so we can query both structures and return the farther of the two answers.

At the general level of recursion, we have an interval $p_i, p_{i+1}, \ldots, p_j$ of points, $i < j$, and a guarantee that any query interval reaching this level of recursion is strictly contained within this interval (excluding both p_i and p_j). At the top level of recursion, $i = 1$ and $j = n$ and we know that the interval contains neither p_1 nor p_n as required. Let $m = \lfloor (i + j)/2 \rfloor$ be the point midway between i and j. We construct a suffix data structure on the left half of points, $p_i, p_{i+1}, \ldots, p_m$, and a prefix data structure on the right half of the points, $p_{m+1}, p_{m+2}, \ldots, p_j$. As above, these data structures can be used to solve any query interval that contains either p_m or p_{m+1} or both (and satisfies the assumption of being strictly contained within the interval $p_i, p_{i+1}, \ldots, p_j$). Then we recursively build data structures in the left half and in the right half for query intervals that do contain neither p_m nor p_{m+1}. In a query, we only need to recurse in one of the halves; we can decide which half overlaps the query interval in constant time by comparing m with the indices of the endpoints of the query interval. In the base case, $j = i$ or $j = i + 1$ and there are no query intervals because of the strict containment, so there is nothing to do.

The recurrence for query time is $T(n) = T(n/2) + O(1) = O(\log n)$. The recurrence for space of the prefix and suffix data structures is $T(n) = 2\,T(n/2) + O(n \log^2 n) = O(n \log^3 n)$. □

Combining Theorems 7 and 8 with Transforms 9 and 10, we obtain the following main result of our paper:

Corollary 11. *There is an $O(n \log^3 n)$-space data structure that supports $O(\log n)$-time halfplane proximity queries on n points in convex position.*

We also mention the implication in the area of dynamic Voronoi diagrams, which follows from combining Theorems 7 and 8 with Transform 9.

Corollary 12. *There is an $O(n)$-space data structure for maintaining a nearest-point or farthest-point Voronoi diagram of a sequence of points in convex position in counterclockwise order. The data structure supports inserting a new point at the end of the sequence, subject to preserving the invariants of convex position and counterclockwise order, in $O(\log n)$ amortized pointer changes per insertion; and supports point-location queries in $O(\log n)$ worst-case time.*

6 Open Problems and Conjectures

Several intriguing open problems remain open. One obvious question is whether the $O(n \log^3 n)$ space of our second data structure can be improved while keeping the optimal $O(\log n)$ query time. One specific conjecture in this direction is the following:

Conjecture 13. *A sequence of n flarb operations, starting from an empty tree, can be implemented at a cost of $O(1)$ amortized pointer changes per flarb.*

We have no reason to believe that our $O(\log n)$ amortized bound is tight. Reducing the bound to $O(1)$ amortized would shave off a $O(\log n)$ factor from our space and preprocessing time. More importantly, it would increase our understanding of dynamic Voronoi diagrams, reducing the $O(\log n)$ amortized update time in Corollary 12 to $O(1)$ amortized. The potential function we use is inherently logarithmic; a completely new idea is needed here for further progress.

On the issue of improving our understanding of dynamic Voronoi diagrams, we pose the following problem:

Open Problem 14. *Is there a data structure for maintaining a Voronoi diagram of a set of points in convex position that allows point to be inserted in $\log^{O(1)} n$ time while supporting $O(\log n)$ point location queries?*

Here we relax the condition that the points be inserted in counterclockwise order, but maintain the restriction that they be in convex position. Although our potential function does not give the result, it is possible that a slight variation of it does.

Finally, it would be interesting to improve the construction time in our second data structure, in particular so that it completely subsumes the first data structure:

Open Problem 15. *Can the pointer changes caused by a flarb be found and implemented in $o(n)$ time, preferably $\log^{O(1)} n$ time?*

We have not been able to fully transform our combinatorial observations about the number of pointer changes into an efficient algorithm, because we lack efficient methods for finding which pointers change. Solving this question would improve our construction time by almost a linear factor, and would provide a reasonably efficient dynamic Voronoi data structure for inserting points in convex position in counterclockwise order.

Acknowledgments

This work was initiated at the Schloss Dagstuhl Seminar 04091 on Data Structures, organized by Susanne Albers, Robert Sedgewick, and Dorothea Wagner, and held February 22–27, 2004 in Germany. This work continued at the Korean Workshop on Computational Geometry and Geometric Networks, organized by Hee-Kap Ahn, Christian Knauer, Chan-Su Shin, Alexander Wolff, and René van Oostrum, and held July 25–30, 2004 at Schloss Dagstuhl in Germany; and at the 2nd Bertinoro Workshop on Algorithms and Data Structures, organized by Andrew Goldberg and Giuseppe Italiano, and held May 29–June 4, 2005 in Italy. We thank the organizers and institutions hosting both workshops for providing a productive research atmosphere. We also thank Alexander Wolff for introducing the problem to us.

References

[AGSS89] Alok Aggarwal, Leonidas J. Guibas, James B. Saxe, and Peter W. Shor. A linear-time algorithm for computing the Voronoi diagram of a convex polygon. *Discrete & Computational Geometry*, 4(6):591–604, 1989.

[Cal05] The Calling. The adventures of Flarb Demingo! `http://www.thecalling.co.za/flarb_pictures.htm`, 2005. See also the fan site, "The Adventures of FLARB with The Calling", `http://www2.fanscape.com/thecalling/streetteam/flarb.html`.

[Cha06] Timothy M. Chan. A dynamic data structure for 3-d convex hulls and 2-d nearest neighbor queries. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms*, 2006. To appear.

[Cro91] Robert G. Cromley. *Digital Cartography*. Prentice Hall, August 1991.

[CT92] Yi-Jen Chiang and Roberto Tamassia. Dynamic algorithms in computational geometry. *Proceedings of the IEEE*, 80(9):1412–1434, 1992.

[dBvKOS99] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer, second edition, 1999.

[Den98] Borden D. Dent. *Cartography: Thematic Map Design*. William C Brown Pub, fifth edition, July 1998.

[DMSW06] Ovidiu Daescu, Ningfang Mi, Chan-Su Shin, and Alexander Wolff. Farthest-point queries with geometric and combinatorial constraints. *Computational Geometry: Theory and Applications*, 2006. To appear.

[DSST89] James R. Driscoll, Neil Sarnak, Daniel D. Sleator, and Robert E. Tarjan. Making data structures persistent. *Journal of Computer and System Sciences*, 38(1):86–124, 1989.

[EGS86] Herbert Edelsbrunner, Leonidas Guibas, and Jorge Stolfi. Optimal point location in a monotone subdivision. *SIAM Journal on Computing*, 15(2):317–340, 1986.

[Epp92] David Eppstein. The farthest point Delaunay triangulation minimizes angles. *Computational Geometry: Theory and Applications*, 1(3):143–148, March 1992.

[MS92] Robert B. McMaster and K. Stuart Shea. *Generalization in Digital Cartography*. Association of American Cartographers, Washington D.C., 1992.

[O'R98] Joseph O'Rourke. *Computational Geometry in C*. Cambridge University Press, second edition, 1998.

[Ove83] Mark H. Overmars. *The Design of Dynamic Data Structures*, volume 156 of *Lecture Notes in Computer Science*. Springer-Verlag, 1983.

[PS93] Franco P. Preparata and Michael Ian Shamos. *Computational Geometry: An Introduction*. Springer, 1993.

[ST83] Daniel D. Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *Journal of Computer and System Sciences*, 26(3):362–391, June 1983.

[ST85] Daniel Dominic Sleator and Robert Endre Tarjan. Self-adjusting binary search trees. *Journal of the ACM*, 32(3):652–686, July 1985.